the Exascale Report™
A Global Perspective
http://TheExascaleReport.com

By Mike Bernhardt
July, 2012

The Intel Perspective on the Company's Acquisition of Cray's Interconnect Technology

An Interview with Intel Vice President and General Manager of Technical Computing, Raj Hazra

A number of industry leaders and analysts have given their opinions on the merits and questions regarding Intel's acquisition of Cray's interconnect hardware assets. And while the attention has been placed on "hardware assets" in these discussions, little emphasis has been given to the human resource assets - the talent and expertise picked up by Intel as part of this deal.

This important transaction took many folks in the HPC community by surprise. Was it a deal that just came out of nowhere? Or was this actually handled with tremendous discretion and secrecy over many months?

As it turns out, this particular deal was handled with great finesse and managed to stay below the radar screens of the HPC rumor mongers.

Intel has always been a difficult company to read when it comes to intentions, direction, and commitment to HPC. After all, job #1 at Intel is microprocessors, and historically HPC has not been able to feed the silicon beast with enough orders for high performance chips to maintain a steady relationship. But it seems like exascale may just have the allure to keep Intel in the game with the same passion we saw in the early 90s.

And, as Intel's Raj Hazra, Vice President of Intel Architecture Group and General Manager of Technical Computing, describes during this Exascale Report feature interview, the human assets acquired from Cray and the knowledge and experience they bring to the table may be even more important to Intel's long term strategy in HPC than the hardware acquisition.

We pick up the conversation with Hazra as he responds to the question of why this was an important transaction for Intel.

[click here to listen to the audio podcast with Intel's Raj Hazra]

HAZRA: So from our perspective, this really is part of our strategy. It wasn't just a flash in the pan. Our strategy is - increasingly as we get toward exascale and then waterfall the technology down into petascale-class machines, but 100x more power efficient, a petascale in a rack for instance, that we would design and integrate the right fabric. You can think of multiple steps we've taken down that path. On top of our Ethernet business we acquired Fulcrum for their technology – their microarchitecture – that could be used in multiple generations of persona. We acquired the Infinband business from Qlogic. And that was for the Infiniband expertise. Clearly even well before we get to exascale, Infiniband is a growing portion of the fabric market. That gave us the competency perspective and market presence. And then the transaction with Cray completed our portfolio of fabric talent as well as a purview of the market from the very high end, into the volume spaces of HPC.

TER: Raj, how does the interconnect technology you acquired from Cray fit in with anything you have been working on internally from the research side in this area?

HAZRA: We haven't and we aren't going to for a little bit, talk about what that fabric or when that integration happens, because that's all part of product plans and we're working through that and we'll use our standard channels to disclose that. We've been looking at and researching what fabrics look like – and what they need to look like in that timeframe to support not only the power, performance, scalability, but also the advent of new programming models like PGAS for instance.

These acquisitions give us the talent as well as the perspective of people who have really done this – and have developed and deployed systems with these high performance fabrics. It gives us the ability to have the right set of people working on them. So the goal is get the talent – get the experience – and we're working through what that roadmap looks like.

TER: Can you tell us how many of the Cray employees made the transition over to Intel?
HAZRA: I don't think we made that public, but it was basically the interconnect core expertise.

[NOTE: In an earlier article with The Exascale Report, Cray CEO Pete Ungaro stated that 74 Cray employees were involved in this acquisition.]

TER: And those folks are all now Intel employees?

HAZRA: They are all Intel employees.

TER: Did most of them have to relocate?

HAZRA: We acquired teams in various locations and they are where they are – today. That doesn't mean that we've quite finished our planning. It's no longer just a fabric effort. As we look at fabrics and their roles going forward with integration, we have to look more holistically at where and how that integration happens with the CPU teams both on Xeon as well as our MIC product line. We'll be optimizing structure. That doesn't necessarily mean moving people all into the same place, but optimizing structure and process to ensure that we have the right mechanisms in place for teams to collaborate and build pieces of silicon that are complex and interdependent. We do that today with processors. We have mutli-site teams working on various aspects of the design and the back end as well. But we have to weave that into the overall fabric – no pun intended.

TER: I think a lot of people have been curious about how the Cray employees might transition into Intel – not logistically – but culturally. Do you have any sense of that at this point?

HAZRA: That's an interesting question and probably best answered by the folks that came in.

Our view always was, and it particularly is – and Mr. Ungaro and I discussed this – we were looking at both synergies of what I would call culture, as well as differences of culture being important. Clearly some synergies around wanting to do cutting edge things: wanting to be on the bleeding edge of technology; the thirst and desire to build something that is stunning; to have flawless execution; to serve customers' needs in the best possible manner. That's a shared

synergy. But we were also looking at, in the spirit of building this capability, at things that we probably didn't know as well and were hoping that some of these acquisitions would teach us. So in many respects, this was not just around the technology of interconnects if you will, but around the whole mindset that comes along with it in terms of understanding system architecture as opposed to component architectures, right down to how do you field large systems with fabrics and processors, and what some of those issues tend to be when you actually put something on the floor and get it up and running. And to the extent that those require new experiences and new competencies as well as new processes and value systems, we appreciated things that are "different" in these acquisitions because we believe that those bring a more complete capability to us, even if they are new to us. So we are valuing not just the synergies and competencies and the value systems and the thirst for cutting edge excellence and technology, but we are also valuing the differences – and the differences actually teach us something because we haven't – in some of these areas – been in the business and make us stronger and more complete as a capability going forward.

TER: That's a very interesting perspective. The acquisition of the human resource assets being such an important strategic part of your think tank –

HAZRA: Yeah, we think it's a great deal for multiple reasons. Not just in the fact that we get great technology and IP, but I continue to say that the great people and the experience they bring in – in areas that were adjacent to us and are now becoming very much a part of our core, is how a company learns and adapts and changes and keeps abreast of where both technology is going – as well as where the user needs are going.

TER: So Raj, let me ask you this. When I think of Cray – I think of a supercomputer company. When I think of Intel, I think of a very large volume commodity processor company. Is there some strategic value of this interconnect technology and those assets to Intel – that goes beyond HPC?

HAZRA: Absolutely. You correctly stated your perspective. I can't argue with your perspective. That just tells us we have work to do because we clearly are increasingly a company that values multiple segments. If you look at the data center business itself, HPC is not a niche business. It's about one third of our data center business. It's one of the faster growing businesses as Diane described at a recent investor day, and it's a growth business for us. Like in any growth business, you invest in it to win. And that's what we are doing in this space.

Supercomputing for us is not a niche for multiple reasons. One is – it's a healthy business in itself. If you look at the growth of Top 500 as kind of a marker for supercomputing – even though supercomputing is broader than the Top 500. The growth in just processors, sockets and flops over the next few years makes just the Top 500 a very, very interesting business in itself. If you look at possibly these over the next year or two of multiple hundred petaflops-plus class machines being built, each of those now represents a single digit percent of the TAM in one machine. And that's just the very top machines. And usually what we've seen is that as the top grows, the bottom continues to grow as well to keep the band roughly consistent over the last 20 years. So, just in terms of sheer volume and revenue in this area, it's interesting today – and it only gets more interesting.

But to the more strategic question you asked, is interconnect an HPC only, even though they're growing, play, the answer is no. The reason is – if you look at multiple segments, whether it's HPC or whether it's what we call high capacity portal datacenters, or you look at very large scale mission critical systems, they are all going through this evolution of a system architecture where power, performance and balance – balance between compute and data movement – are becoming important. And while the fabrics or interconnects may actually be different across multiple segments, the role of the interconnect in being the central point of the system architecture and defining the difference between a good and a badly balanced system is true for all of those segments. And therefore, the fabric, if you look at the interconnect for anything more than a single node system over time is becoming a fundamental value proposition of the data center architecture. We're in all aspects of the data center business. The compute platform, the networking platform, the storage platform, and so we see this as almost integral to our portfolio – today, and even more critical in the future as integration drives more and more of that capability into the pieces of silicon that we build, and potentially extend our business through software on top of that.

TER: So there was a statement made recently – someone had stated that they didn't believe Intel was truly interested in exascale because they didn't perceive exascale as offering nearly enough volume of processors to be pushed out for a company like Intel. But you are already saying that HPC in itself is starting to recognize tremendous volume of processors, therefore it's becoming an even more attractive business for Intel – so let's talk about exascale. Obviously still a long way off and possibly creeping out a little bit farther and farther each year. Interconnect a huge part of that. There's probably no idea on the drawing board right now as to what processor Intel will be putting out at that point in time -

HAZRA: It's on our drawing board.

TER: So, can you give me a picture of how you see things changing in terms of the fabric – the landscape –of what we are going to be looking at for exascale-class systems?

HAZRA: You asked a very interesting set of interrelated questions Mike. You started by asking why are we interested in exascale. And I think our interest in exascale isn't - personally, I'm not as excited about a single exascale machine at 20MW. Of course, to the extent that such is built, and we help power it – power it as provide the critical technologies to it and it serves interesting missions, I'm all for it.

The real value of exascale, and the real value of petascale and terascale in previous years, has been that it fundamentally changes the ecosystem in a set of ways. And the best way to state this is – when we reach exascale at 20MW – you're right – if you look at petascale today, there are probably going to be 4 or 5 or 6 computers in the world initially at exascale. And then over time, a lot of the Top 500 will get there. But that's what's happening with petascale today. The question isn't just the business value of what 4, 5, 6, or 10 machines at exascale at 20MW – or near exascale – the real value if you can get there is that you have now provided a capability that can fundamentally change the performance per watt envelope of compute to be able to serve petaflops in a single rack.

So, synonymous with reaching an exascale system at 20MW, our goal is to enable a petaflop-class machine, which would get you in the Top Ten today, in a rack. Now a rack is typically a work group class HPC machine today. Think of the possibilities of being able to field, in a rack, a petaflop of compute. The TAM it opens up is huge. So to us, the value of exascale is you pull the top up in capability and by doing so you actually benefit the entire pyramid of HPC so that it actually grows the market. You have many more new entrants who can now use this capability because you have now changed the performance per watt, and therefore parts of the TCO equation very significantly. So I sometimes stand up and say, I'm very excited about the vision of a very large data center with a beautifully humming exascale machine at 20MW, but I'm probably a 100x more excited by millions of racks that are only consuming 20 to 30 kilowatts, but each fielding a petaflop of compute. That is the promise of exascale. That's what petascale has done. You know, there is probably, if I count on the Top 500, maybe 7-10 petaflop machines today and there may be a few more in a couple of weeks [referring to ISC12], but look at what it's done to the rest of the pyramid. Petascale at maybe 6, 7, 8 or 10 megawatts allowed terascale machines to be built at hundreds of kilowatts. Just imagine if you could build an exaflop at 20MW. You could build thousands or hundreds of thousands of petascale machines at what today is half a rack. The market expansion of HPC when you democratize compute capability that way is stunning. That's what excites us about the marketplace.

TER: Are there any other initiatives coming from Intel in terms of building community as we go down the path toward exascale?

HAZRA: Yes, there certainly is. It's almost a question of where do I start and where do I end – and how much time do you have. The one thing that is very clear to us is that exascale, even though it sounds so system and hardware centric , for us is an ecosystem statement. Because just as important as the question of how do you build an exascale machine, it is equally important to answer the companion question of what are you going to do with an exascale machine? And it has to do with things like what applications are going to be available and able to scale and take advantage, what programming models will people have to learn to write those applications, and in order to do that, we've got a plethora of activities that are adjacent to but complimentary to our core business of building hardware and software tools. For example, around our exascale labs in Europe, one of the labs is looking at space weather as a potential application. One could argue about the value of space weather and when my nightly news is going to start talking about the weather in space, but it is fundamental science and it is an extremely good example of an application that needs exascale-class computing, and it furthers our frontiers of all those things. How do you build these applications? How do you understand fundamental algorithms and their interactions? How do you program these things at large scale? So we have efforts on applications through our labs in Europe. We have at the Intel Research Labs level interesting collaborations going on in parallel computing. And now we have just announced an effort in Big Data because we are starting to see Big Data as both an interesting complex set of workloads for HPC but in many ways also in forming the evolution of system architectures as we move forward. So we have community efforts along with universities and other collaborators on understanding these frontiers of next generation workloads. And then of course we have our next area of work in the community which is just bringing people to

use more and more of HPC. We typically call this the missing middle. And apropos to my comments earlier about a petaflop in a rack, the first thing you would ask is who would use a petaflop in a rack? In the manufacturing space and in the life sciences space, there is a number of people today who we can make a case would economically benefit in their mission from HPC but aren't using it for a multitude of reasons. So how do you get them over that barrier to use? And that doesn't have to wait for exascale. That's starting at the scale of computing we have today. But we hope that continues to build a more vibrant community as the compute capability continues to move up toward exascale.

TER: Last question. So it's been about two years since the ISC when Intel put a stake in the ground saying you will get us to exascale by 2018 and at 20MW. How does that look today?

HAZRA: What we said is – we see a technical path to getting there. It has some large challenges in terms of power, performance, resiliency, scalable programming models, and if there is conviction in the industry and conviction amongst governments to actually go fund this R&D, that we would do our part but we could see a technical path to getting there.
Now we remain committed to that statement. But we said we can't do it and won't do it alone. To the extent that there is the partnership and the skin in the game from the right places, and the collaboration, we think it's still very much in the cards and we're hoping that as we do our part – and others do their part – that we will get there.

< < < < < > > > > >